www. chameleoncloud.org

From Cloud to Edge: Building Instruments for Today's Science

Kate Keahey

University of Chicago / Argonne National Laboratory

keahey@uchicago.edu

NCAR Seminar, 01/23/25

TESTBEDS AS SCIENTIFIC INSTRUMENTS

- Research Infrastructure: a tool for computational experimentation
  - The experiments we can think about are unlimited…
  - …but in practice we can carry out only those that are supported by an instrument that allows us to deploy, capture (observe and measure), and record relevant scientific information
- Of Telescopes and Tokamaks
  - Exploratory instruments (tokamaks): deploy then measure -- Chameleon
  - Discovery/observational instruments (telescopes): measure – FLOTO infrastructure
- Challenges in scaling and adaptability

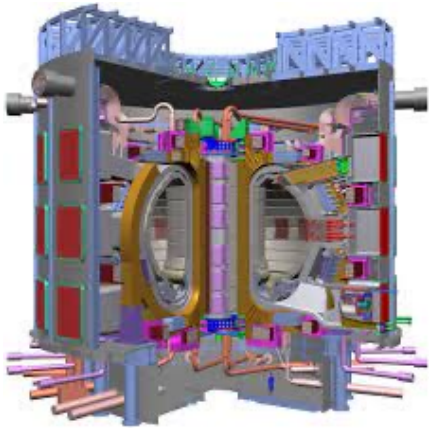

CHAMELEON: AN EDGE TO CLOUD TESTBED

- Chameleons like to change – testbed that adapts to your experimental needs
  - **From bare metal reconfigurability/isolation** -- KVM cloud – to containers for edge (**CHI@Edge**)
  - Capabilities: power on/off, reboot, custom kernel boot, serial console access, etc.
- From large to small – diversity and scale in hardware:
  - **Supercomputing datacenters** (UC/ALCF, TACC, NCAR) over 100G network – to **edge devices**
  - **Diverse:** FPGAs, GPUs, NVMe, NVDIMMs, Corsa switches, edge devices via CHI@Edge, etc.
  - **Distributed: CHI-in-a-Box** sites at **Northwestern and UIC** – and now also **NRP**!
- Based on mainstream open source – proud to be cheap!
  - 50% leveraging and influencing **OpenStack** + 50% "special sauce" (incl. fed id)
- Promoting digital artifact sharing
  - Integration with **Jupyter** for non-transactional experiment packaging
  - **Trovi** for experiment sharing and discovery, **Chameleon Daypass** for access sharing
  - Reproducibility and education: digital sharing killer apps!


Chameleon
700+
Papers
published
1,000+
Unique
projects
11,000+
Users
9

CHAMELEON HARDWARE

Coming soon: Dell XE9640,
2x Intel 9468 CPU / 4x Nvidia H100

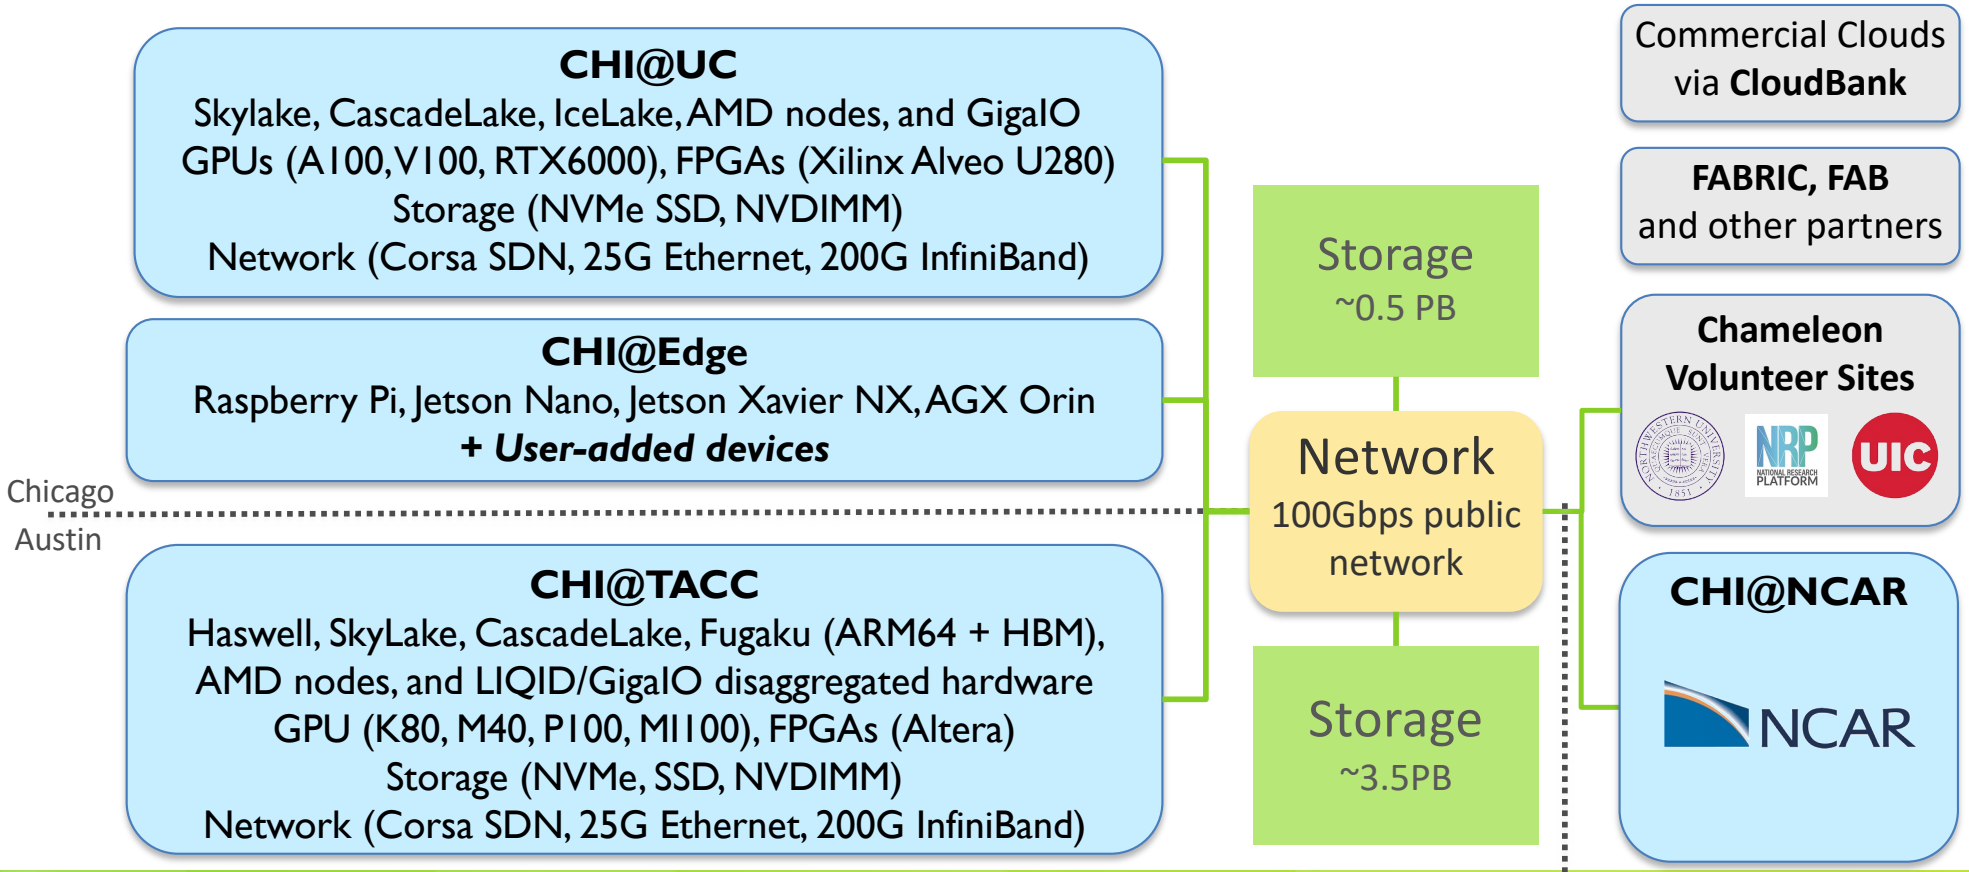

CONFIGURATION HIGHLIGHT: CHI-IN-A-BOX

- CHI-in-a-box: packaging of CHameleon Infrastructure (CHI)
  - Internal packaging of a commodity-based testbed
  - Packages the system as well as the operations model
  - Hub and spoke management, version-controlled site configuration management as code, containerization, monitoring, detection, and remediation tools
  - Support for Bring Your Own Device (BYOD) model: Doni allows administrators to dynamically enroll resources, define availability windows, and streamline operations
- Deployment
  - Deployed Associate/Volunteer Sites: NCAR, Northwestern, NRP, and UIC
  - Independent testbed: ARA
  - In conversation/progress: OCT/U Mass, FIU, ORNL, KTH (edge/wireless only), NUS, and others

Paper: "CHI-in-a-Box: Reducing Operational Costs of Research Testbeds ", PEARC'22

EXPERIMENT STRUCTURE

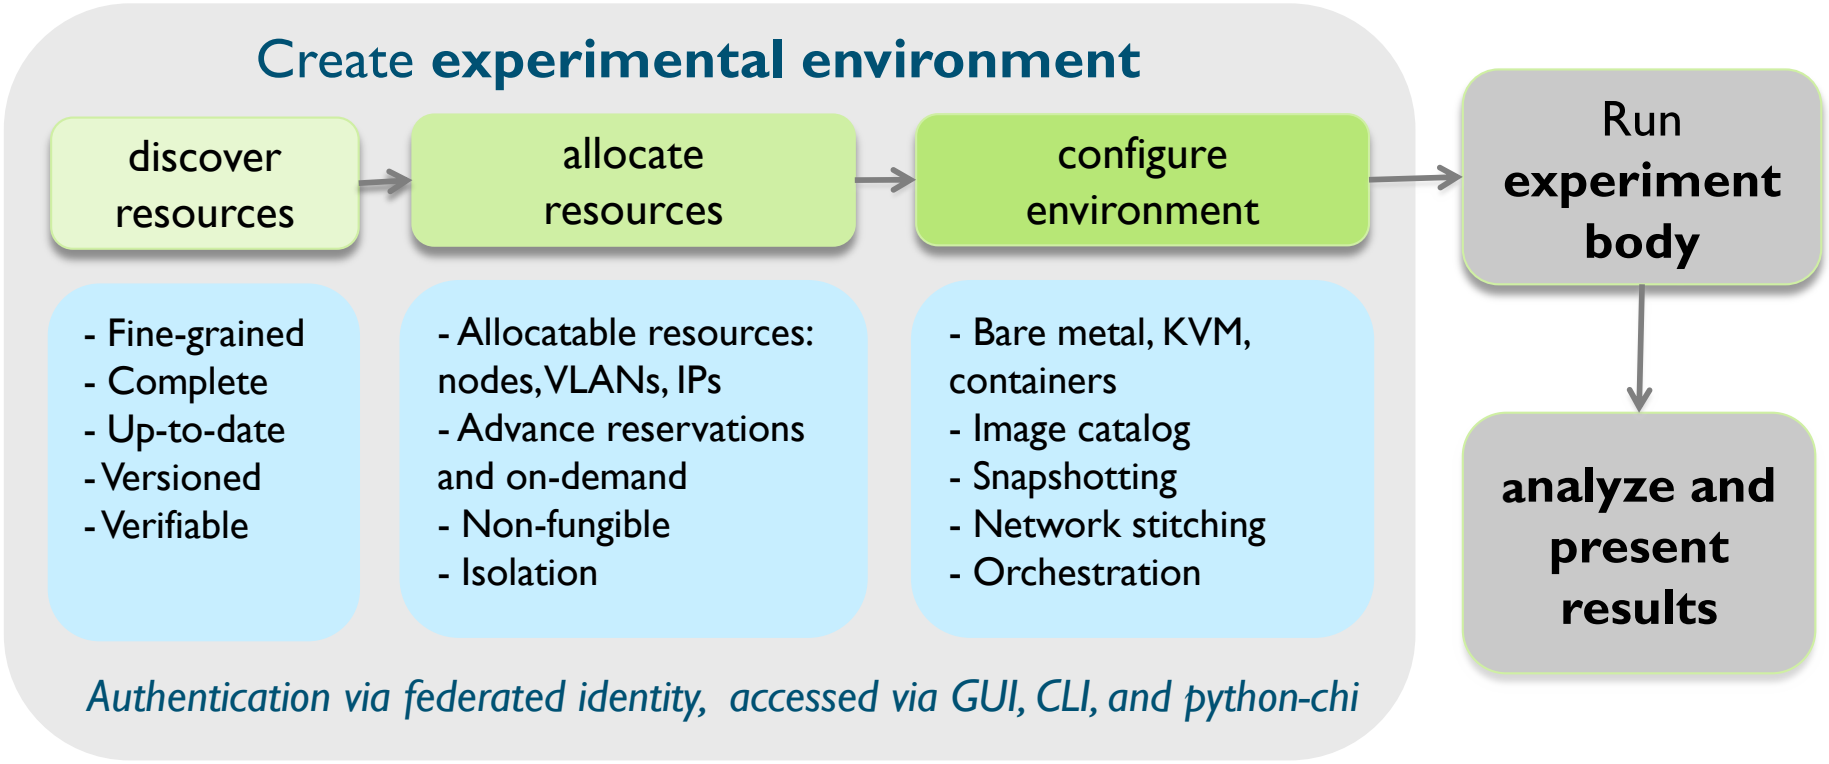

Paper: "Lessons Learned from the Chameleon Testbed", USENIX ATC 2020

NOT JUST A TESTBED, A COMMUNITY

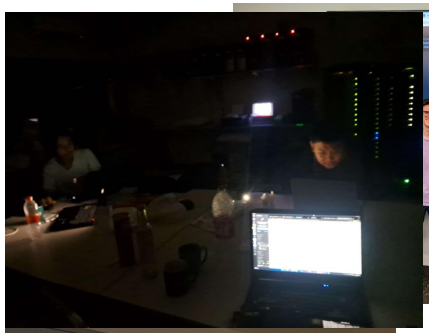

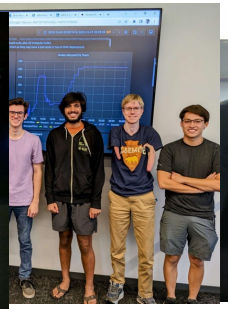

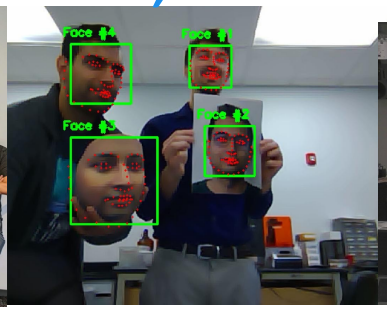

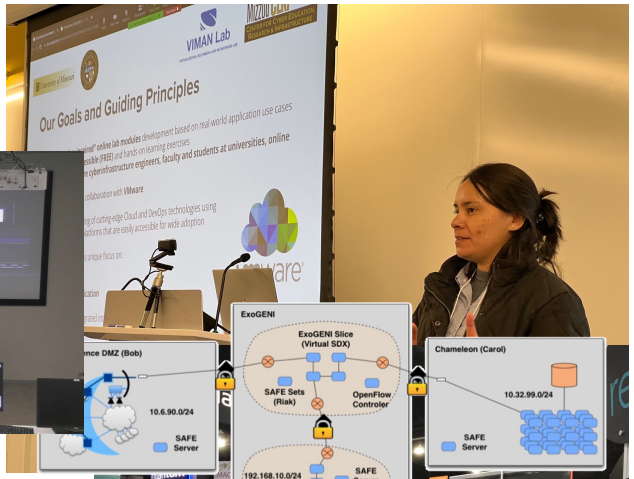

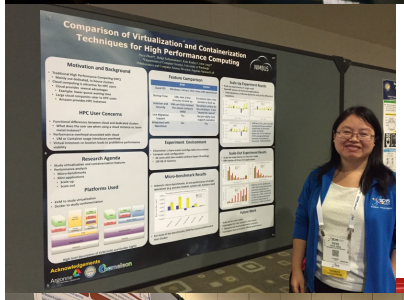

Supporting research projects in architecture, operating systems design, virtualization, power management, real-time analysis, security, storage systems, databases, networking, machine learning, neural networks, data science, and many others.

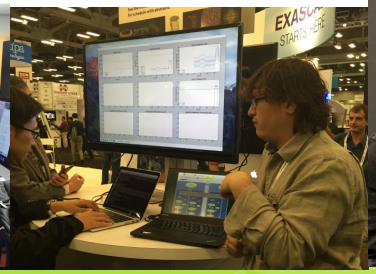

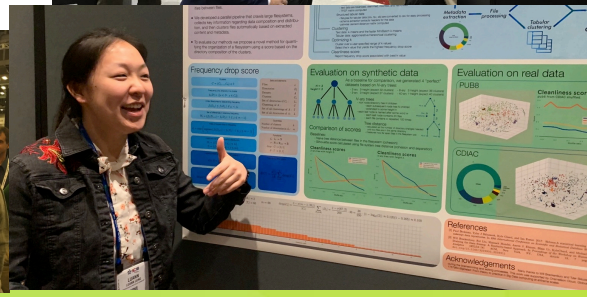

Chameleon www.chameleoncloud.org

Check out user experiment stories on our blog:
https://www.chameleoncloud.org/blog/category/user-experiments/

PRACTICAL REPRODUCIBILITY

Practical reproducibility == feasible enough to be a mainstream method of scientific exploration

- Can digital experiments be as sharable as papers are today?
- Is there a library I can go to and find experiments to play with?
- Can I simply integrate somebody's model into my research instead of reinventing the wheel and get to a new result faster?
- Can I discover something new through playing with somebody else's experiment?
- Can I develop exercises for my class based on most recent research results?

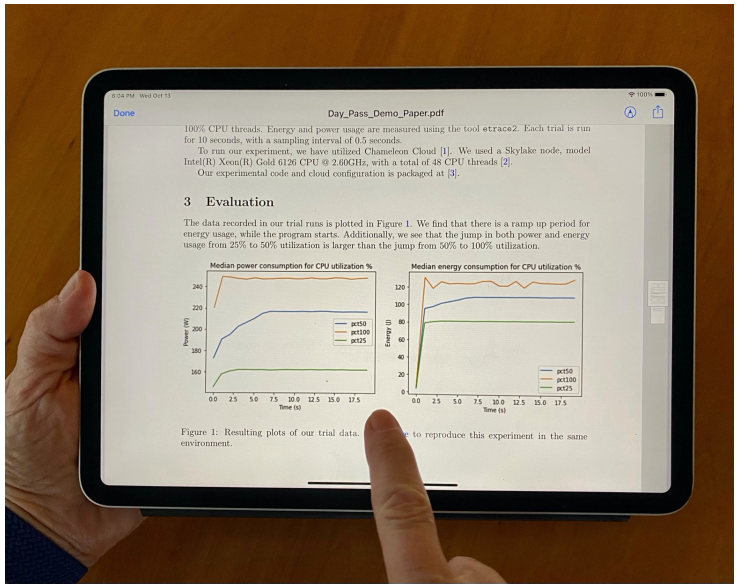

https://repeto.cs.uchicago.edu

WHAT DO WE HAVE?

- **Open platforms** are essential for sharing – especially in computer science
  - Open, version-controlled hardware
  - Non-fungible resources
- **Experimental environment setup**
  - Disk images, orchestration templates, and other artifacts
  - Thousands of images, orchestration templates, digital artifacts of various kinds
- Are we there yet?

A car without a road

Paper: "The Silver Lining", IEEE Internet Computing 2020

WHAT IS MISSING?

- **Shared hardware**: Chameleon daypass for open access
- End-to-end **packaging** with literate programming
  - Credential integrated JupyterLab environment: convenience of notebook + power of testbed
  - Imperative, non-transactional, annotated
- Trovi: an experiment sharing **repository**
  - Portal to present, browse, filter, and find
  - Integrated with Jupyter/Chameleon, Swift, Zenodo, and github
  - Open APIs: FABRIC, Jetstream2, and others

Practical reproducibility:

cost-effective enough to be mainstream

Effective packaging via a "compute capsule"

Finding and sharing experiments integrated with platform

Open, integrated access to all aspects of experiment

Paper: "Three Pillars of Reproducibility", ReWords'23

TROVI ARTIFACT GROWTH

Total artifacts	187
Unique authors	81
Daypass enabled artifacts	22
Chameleon supported	15
Repeto badges	42
Fount badges	24
Max access count ("launches")	1256
Mean access count	50.32
Max unique access count	181
Mean unique access count	11.64
Max unique cell execution	148

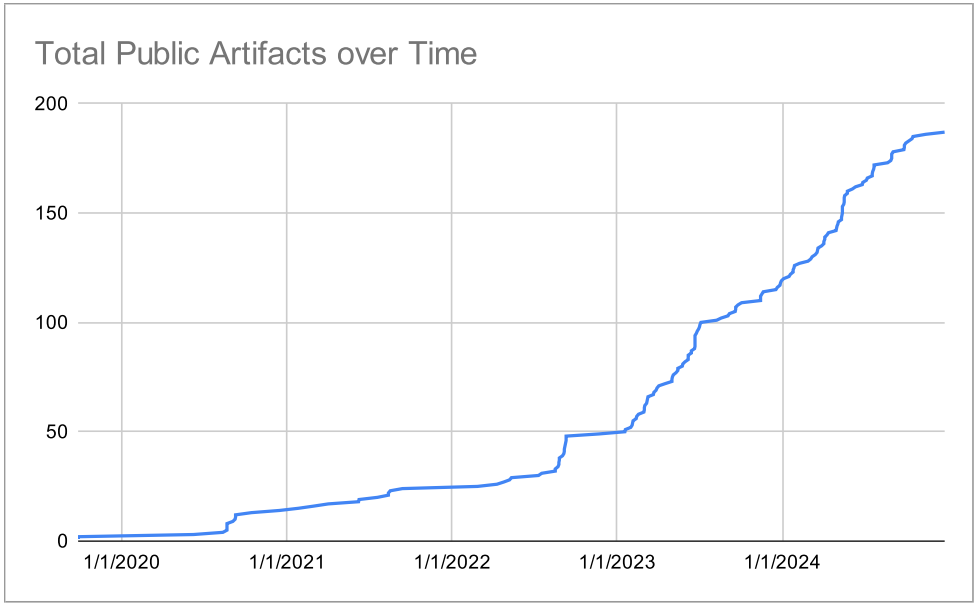

FROM CLOUD TO EDGE WITH CHAMELEON

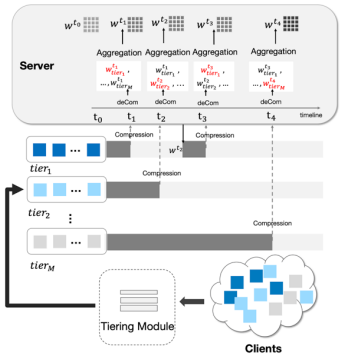

federated learning

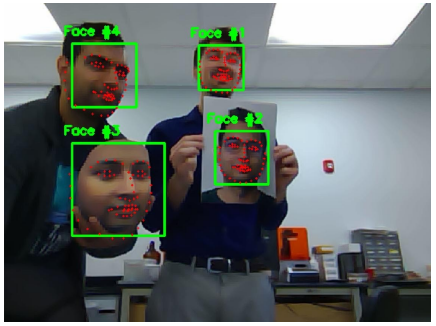

biometrics

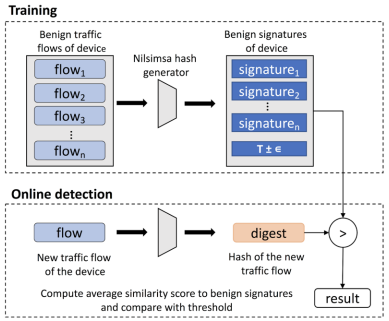

network traffic fingerprinting for IoT devices

- Increasingly more Chameleon project applications working on IoT/edge
- Simulation/emulation don't always provide the answer: What are the impacts of this approach on power management on edge device? How will the performance transfer to edge? Can we measure the impact of distribution/networking for edge/cloud applications?
- **Goal: "realistic edge to cloud experiments from one Jupyter notebook"**

HIGHLIGHT: CHI@EDGE

A lot like a cloud!
All the features we know and love – but for edge!
"Edge to cloud from one Jupyter notebook."

Not at all like a cloud!
Location, location, location!
IoT: cameras, actuators, SDRs!
Not server-class!
And many other challenges!

- CHI@Edge: all the features you love in CHI, plus:
  - Reconfiguration through non-prescriptive **container deployment** via OpenStack interfaces (using K3 under the covers)
  - Support for "standard" **IoT peripherals** (camera, GPIO, serial, etc.) + easy for you to add support for your own peripherals
  - **Bring Your Own Device (BYOD): Mixed ownership** model via an SDK with devices, virtual site, and **restricted sharing** – building on OpenBalena

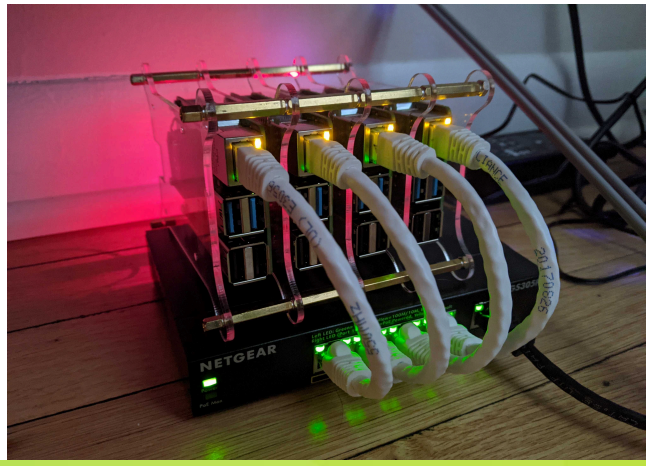

Paper: "Chameleon@Edge Community Workshop Report", 2021

FROM EXPLORATION TO OBSERVATION: THE FLOTO PROJECT CASE STUDY

- Why broadband monitoring?
  - Technical questions: what happens in conditions of oversubscription?
  - Policy questions: is there a "digital divide" in our society?
  - Modeling questions: what assumptions about broadband are realistic?
- Measuring broadband – different approaches/applications depending on context, objective, use case, etc.
  - Netrics: open-source library of standard network diagnostic tools (ndt7, speedtest, ping, traceroute, etc.) for continuous, longitudinal network measurement
  - Others: e.g., residential versus rural broadband and other use cases
- **Can we use CHI@Edge as a large observatory instrument for broadband monitoring?**
- **Approach:** connect a "measurement box" to the router and run tests
- Collaboration with Nick Feamster & his UChicago team

THE DEVICES

- Raspberry Pi 4 (8GB)
- Additional Components
  - MicroSD Cards (32GB)
  - CAT 6 Ethernet Cable
  - Power Cord
- Optional: PoE+ HATs to enable deployment in locations with scarce power sources
- Inventory:
  - 1,000 devices finished arriving at the end of June'23
  - In the process of deploying them for broadband research

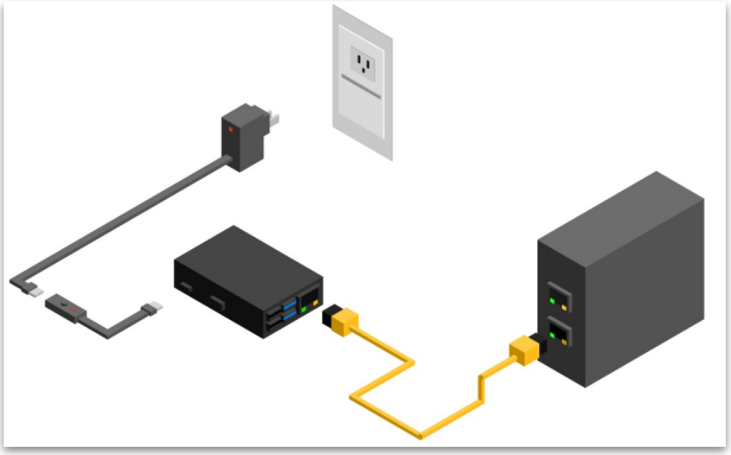

DEVICE MANAGEMENT LAYER

- Onboard, offboard, and repurpose devices
- Devices self-enroll
  - 0 touch device enrollment (after imaging)
  - Alternatively, flash with our image to enroll your own device
- Configuration management
  - Update and deploy without physical access, stateless operating system, includes software and device configuration, can be pinned to releases
- Robust remote management features
  - View status and statistics, create and manage deployments, trigger appropriate actions (e.g., send mail)

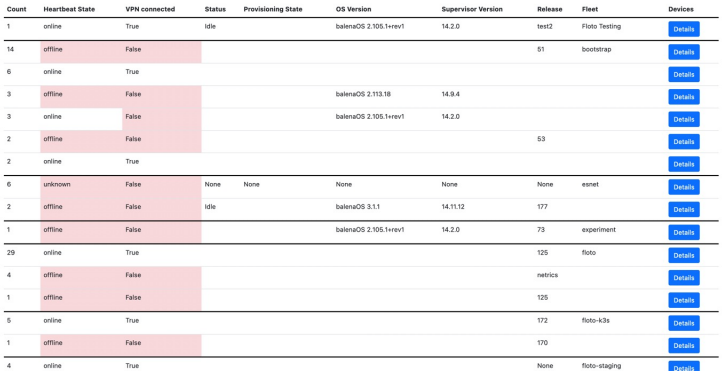

Count	Heartbeat State	VPN connected	Status	Provisioning State	OS Version	Supervisor Version	Release	Fleet	Devices
1	online	True	Idle		balenaOS 2.105.1+rev1	14.2.0	test2	Floto Testing	Details
14	offline	False					51	bootstrap	Details
6	online	True							Details
3	offline	False			balenaOS 2.113.18	14.9.4			Details
3	online	False			balenaOS 2.105.1+rev1	14.2.0			Details
2	offline	False					53		Details
2	online	True							Details
6	unknown	False	None	None	None	None	None	esnet	Details
2	offline	False	Idle		balenaOS 3.1.1	14.11.12	177		Details
1	offline	False			balenaOS 2.105.1+rev1	14.2.0	73	experiment	Details
29	online	True					125	floto	Details
4	offline	False					netrics		Details
1	offline	False					125		Details
5	online	True					172	floto-k3s	Details
1	offline	False					170		Details
4	online	True					None	floto-staging	Details

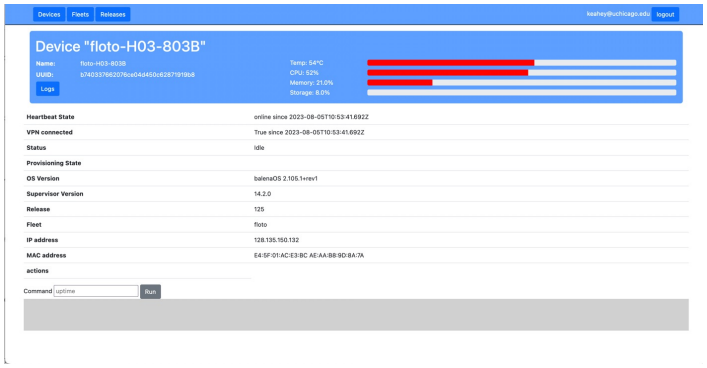

DEVICE MANAGEMENT LAYER IMPLEMENTATION

- Federated Identity login (via GlobusAuth)
- Based on openBalena (open source project underlying BalenaCloud platform)
  - Extensions include support for federated identity, managing device collections, ability to execute ad-hoc shell commands on device system containers, and others
  - Dashboard to expose management functions
- Control and data services deployed on HA infrastructure
- Devices run minimal, stateless, vetted images with the application management layer software
- Access via the dashboard or a CLI

APPLICATION MANAGEMENT LAYER

- Supports deployment of applications on device fleets via a system container
- Applications are packaged as Docker containers
- Users can reserve overlapping or non-overlapping timeslots for application deployment so as not to conflict with other deployments
- Generic data streaming implemented as a “system application”
- Multi-container applications deployed via docker-compose syntax

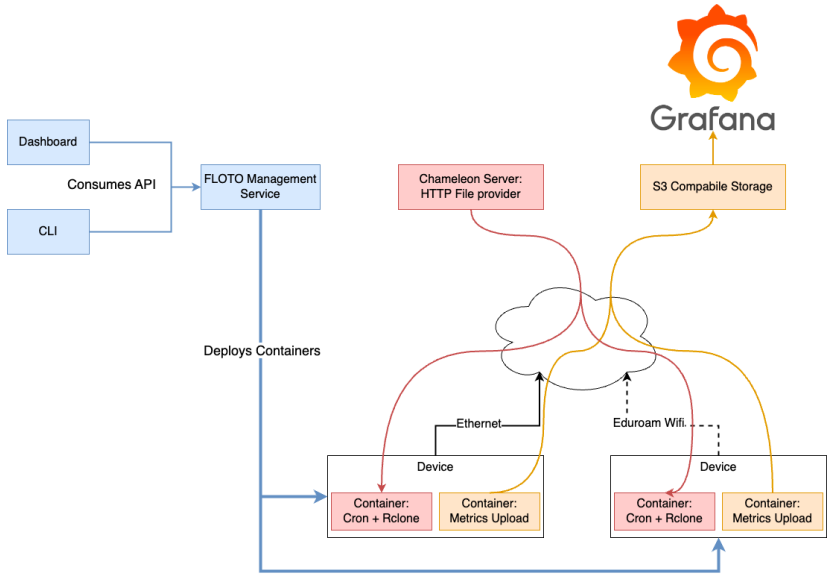

DATA COLLECTION AND ANALYTICS

- Software on devices uploads results directly to cloud storage
- Data curation and analysis pipelines process raw data for investigation
- Raw and processed data available for use (after anonymization)
- Current, small scale, using AWS Lambda for analysis
- After scale up, use Apache Kafka as message broker, data consumers will subscribe to real time topics for up to date results
- Grafana dashboards with time-series visualizations to monitor data in near real-time

Paper: "Discovery Testbed: An Observational Instrument for Broadband Research", eScience'23

INSTRUMENT ADAPTABILITY

- **What knobs can I turn on this instrument?**
- Deployment scope: deploy the devices in a different area
- Hardware: combine devices with different IoT peripherals (e.g., GPS)
- Application: adapting “sensing abilities” programmatically
- Data aggregation: different methods for different applications
- Data: ask different questions of the data

FLOTO: GIVING BROADBAND MEASUREMENT AN EDGE

- Scientific instrument for measuring broadband
- Deploy 1,000 Pis nationwide (~500 so far)
  - Chicago, IL; Milwaukee, WI; San Rafael, CA
  - Marion County, IL; Beaver Island, MI -- and others
- Measurement Applications
  - Netrics; Measurement Lab's (MLab) Measurement Swiss Army Knife (MSAK) toolkit; RADAR toolkit; NetUnicorn; rural broadband tests (ARA) – and others
- Data
  - 13M data points, spanning 17 providers (national and local), across multiple different technologies
  - Publicly available on FLOTO website
- How powerful is this dataset?
  - Marion County: 32% of sampled households below the federal threshold
  - Beaver Island: area challenge to FCC -> reassessment of broadband coverage

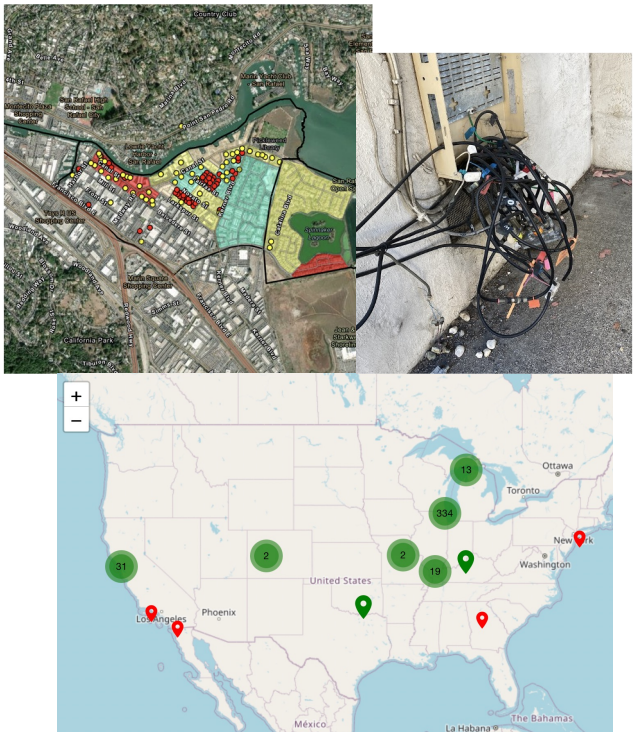

floto.cs.uchicago.edu

MEASURING RURAL WIRELESS

- Collaboration with ARA project
- Assessing the quality of rural 5G networks
  - Measuring device to device latency
  - Clock synchronization
  - Comparing over different network fabrics
- Deployed 6 Raspberry Pi devices with 5G connectivity in rural Iowa
- Latency measurements: GPS-based time synchronization for precise measurements (4000x more precise than NTP over 5G)
- Tested using Hadoop
- Hey presto: 5G networks can support distributed computing with performance comparable to wired connections!

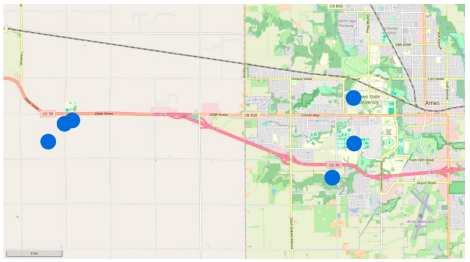

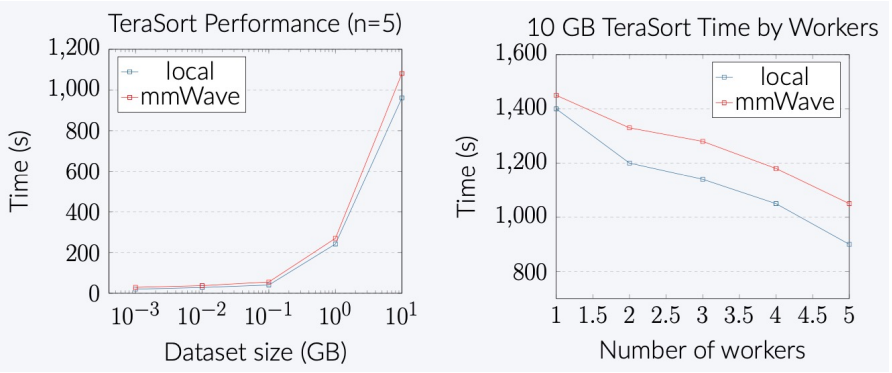

Zack Murry, University of Missouri

NCAR WEATHER SENSING STATIONS

- openIoTwx: NCAR 3D printed weather stations
- Richer continuum: IBIS SBCs connecting to openIoTws via LoRa
  - Exploring power (4x factor), connectivity (cellular vs aggregation via LoRa), sensing (additional camera sensors), and processing (to e.g., reduce size of data) trade-offs
- Future challenges
  - Image-based weather prediction methods, scaling up to create dense, high-resolution weather monitoring networks, and assessing long-term reliability in diverse outdoor environments

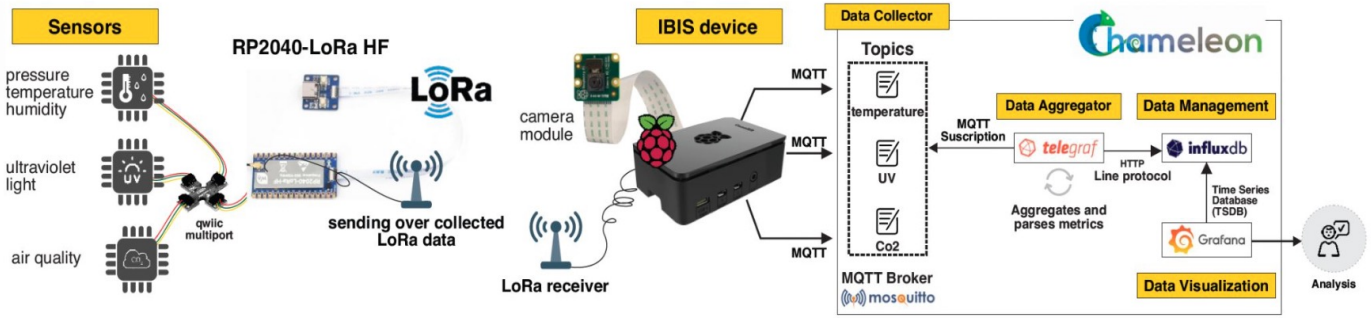

William Fowler, Tufts University

SENSOR STATIONS FOR MARINE AND COASTAL ECOSYSTEMS

- Smart buoy system: sensor stations for oceanic data collection (water quality, water movement, water levels, etc.)
- Collaboration with FIU
- Integrated multiple environmental sensors with IBIS infrastructure
- Demo deployment with real and simulated data
- Implemented cloud-based data visualization system
- Collaboration with FIU

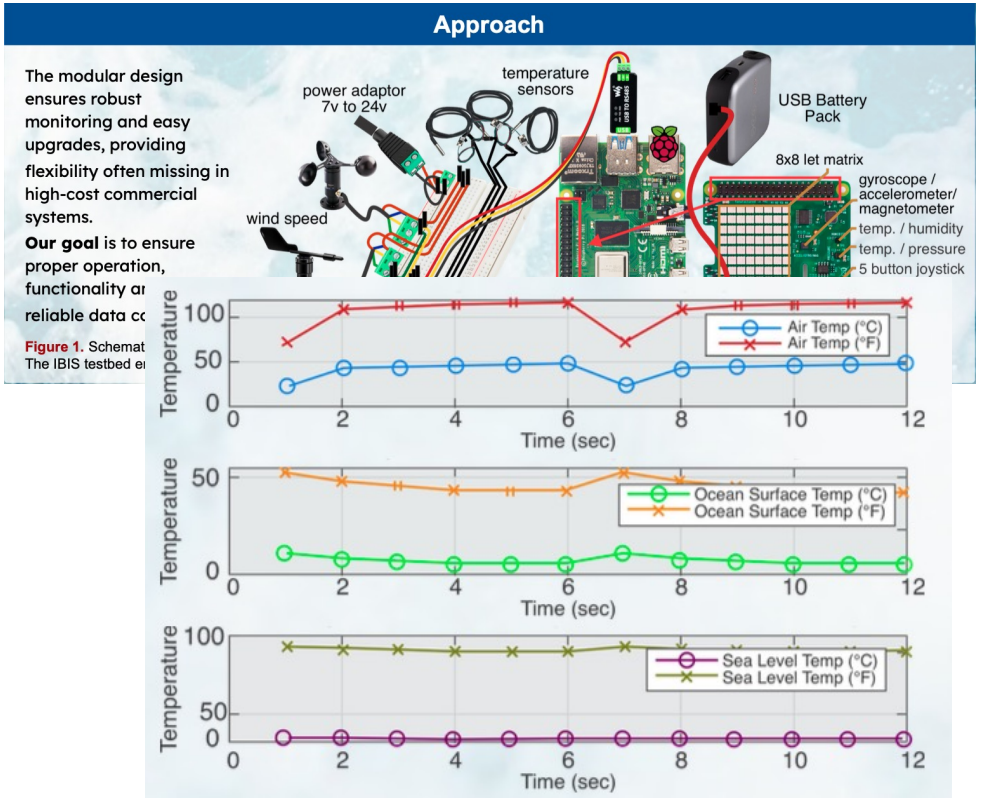

AUTOLEARN: A CASE STUDY

A collection of courselets exploring concepts in autonomous driving

- **Contains three types of courselet layers:**
  - Data collection (actual car versus simulator)
  - Machine Learning courselets training models
  - Verification via self-driving (actual car versus simulator)
- **Supports different emphasis and different pathways through the curriculum:**
  - Introduction to engineering might emphasize driving the actual car
  - Machine learning focus might use the simulator
- **Contain suggestions for exercises and individual exploration:**
  - E.g., digital twin combining simulator and experimental driving

Paper: "AutoLearn: Learning in the Edge to Cloud Continuum", EduHPC'23

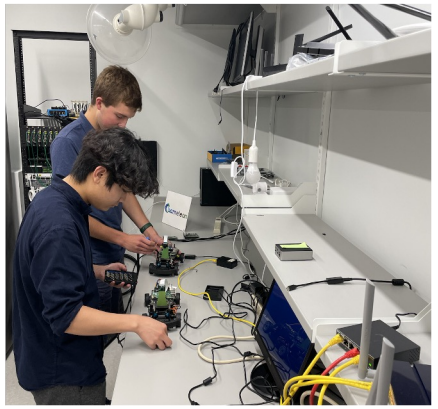

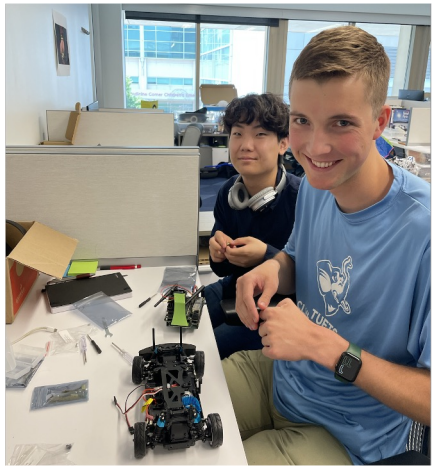

REU 2023 students working on hardware setup for autonomous vehicles

AND OTHERS…

- Predicting air quality with federated learning
- Soundscaping and forestry data analysis
- Precision agriculture: optimizing greenhouse environments
- Meteorologic monitoring system for ML-based weather forecasts
- And more…

FOR BETTER OR WORSE, SCIENTIFIC INSTRUMENTS SHAPE A FIELD

research highlights

DOI:10.1145/2209249.2209271

Check for updates

Technical Perspective
For Better or Worse, Benchmarks Shape a Field

By David Patterson

LIKE OTHER IT fields, computer architects initially reported incomparable results. We quickly saw the folly of this approach. We then went through a sequence of performance metrics, a victim of its own success. The SPEC organization has been selecting old programs written in old languages that reflect the state of programming in the 1980s. Given the 1,000,000X improve-

Given this measurement framework, the authors then measured eight very different Intel microprocessors built over a seven-year period. The authors evaluate these eight micropro-

We're all snow plough drivers now!

We're here to change

www.chameleoncloud.org